

# Efficient Computation of Gaussian Likelihoods for Stationary Markov Random Field Models

Joseph Guinness <sup>\*1</sup> and Ilse C. F. Ipsen <sup>†2</sup>

<sup>1</sup>*North Carolina State University, Department of Statistics*

<sup>2</sup>*North Carolina State University, Department of Mathematics*

June 2, 2015

## Abstract

We introduce methods for efficiently computing the Gaussian likelihood for stationary Markov random field models, when the data locations fall on a possibly incomplete regular grid. The calculations rely on the availability of the covariances, which we show can be computed to any user-specified accuracy with fast Fourier transform algorithms. Several methods are presented, covering models with and without additive error terms and situations where either conserving memory or reducing computation time are favored, and some of the algorithms are easily parallelized. The examples presented highlight frequentist inference, but access to the likelihood allows for Bayesian inference as well. We demonstrate our results in simulation and timing studies and with an application to gridded satellite data, where we use the likelihood both for parameter estimation and likelihood ratio model comparison. In the data analysis, stochastic partial differential equation approximations are outperformed by an independent block approximation.

## 1 Introduction

In a wide range of scientific and engineering applications, such as imaging, agricultural field trials, climate modeling, and remote sensing, spatial data are observed or reported on regular grids of locations, and the Gaussian process model is commonly used to model the data directly, or indirectly as a stage in a hierarchical model (Rue et al., 2009; Banerjee et al., 2014). Let  $Y(\mathbf{x}) \in \mathbb{R}$ ,  $\mathbf{x} \in \mathbb{Z}^2$  denote a real random variable indexed by the two-dimensional integer lattice. We say that  $Y(\cdot)$  is a Gaussian process if, for any  $n \in \mathbb{N}$  and any set of locations  $\mathbf{x}_{1:n} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ , the vector  $\mathbf{Y} = (Y(\mathbf{x}_1), \dots, Y(\mathbf{x}_n))'$

---

<sup>\*</sup>jsguinne@ncsu.edu

<sup>†</sup>ipsen@ncsu.edu

has a multivariate normal distribution. A common decomposition of a Gaussian process is

$$Y(\mathbf{x}) = \mu_\beta(\mathbf{x}) + Z(\mathbf{x}) + \varepsilon(\mathbf{x}),$$

where  $\mu_\beta(\mathbf{x})$  is a nonrandom function depending on mean parameter  $\beta$ ,  $Z(\cdot)$  is a Gaussian process with mean  $E(Z(\mathbf{x})) = 0$  and covariance function  $\text{Cov}(Z(\mathbf{x}), Z(\mathbf{y})) = K_\theta(\mathbf{x}, \mathbf{y})$  depending on covariance parameter  $\theta$ , and  $\varepsilon(\mathbf{x}) \sim \text{i.i.d. } N(0, \sigma^2)$ . The usual terminology in this scenario is to call  $Z(\cdot)$  a latent process, since it is not directly observed, and to say that the model contains a nugget effect, referring to the presence of the  $\varepsilon(\mathbf{x})$  term, which is used to model microscale variation or measurement error.

We define the mean vector  $\boldsymbol{\mu}_\beta := E(\mathbf{Y}) = (\mu_\beta(\mathbf{x}_1), \dots, \mu_\beta(\mathbf{x}_n))'$ , the random vector  $\mathbf{Z} = (Z(\mathbf{x}_1), \dots, Z(\mathbf{x}_n))'$  and its covariance matrix  $\Sigma_\theta$ . Then the covariance matrix for  $\mathbf{Y}$  is  $\Sigma_\theta + \sigma^2 I$ , where  $I$  is an identity matrix of appropriate size. The parameters can be estimated using the loglikelihood function for  $\beta$ ,  $\theta$ , and  $\sigma^2$  from  $\mathbf{Y}$

$$L(\beta, \theta, \sigma^2) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log \det(\Sigma_\theta + \sigma^2 I) - \frac{1}{2} (\mathbf{Y} - \boldsymbol{\mu}_\beta)' (\Sigma_\theta + \sigma^2 I)^{-1} (\mathbf{Y} - \boldsymbol{\mu}_\beta).$$

Memory and time constraints begin to prohibit storing and factoring  $\Sigma_\theta + \sigma^2 I$  when  $n$  is between  $10^4$  and  $10^5$ . Since the inclusion of a non-zero mean vector does not introduce a substantial additional computational burden, we assume  $\boldsymbol{\mu}_\beta = 0$  in our discussion of the methods, to simplify the formulas and focus our attention on estimating covariance and nugget parameters.

There is significant interest in identifying flexible covariance functions  $K_\theta(\mathbf{x}, \mathbf{y})$  for which the likelihood function—or an approximation to it—can be computed efficiently (Sun et al., 2012). Gaussian Markov random field (GMRF) models (Rue and Held, 2005) are of particular value because they provide a flexible class of models that induce sparsity in  $\Sigma_\theta^{-1}$ , a feature that can be exploited by sparse matrix factorization algorithms. A GMRF on the infinite integer lattice  $\mathbb{Z}^2$  is naturally specified by the conditional expectation and variance of each observation given every other observation on  $\mathbb{Z}^2$ . Formally, we write

$$Z(\mathbf{x}) | \{Z(\mathbf{y}) : \mathbf{y} \in \mathbb{Z}^2 \setminus \mathbf{x}\} \sim N\left(- \sum_{\mathbf{y} \in \mathbb{Z}^2 \setminus \mathbf{x}} \frac{\theta(\mathbf{x}, \mathbf{y})}{\theta(\mathbf{x}, \mathbf{x})} Z(\mathbf{y}), 1/\theta(\mathbf{x}, \mathbf{x})\right), \quad (1)$$

where  $\theta(\cdot, \cdot)$  is a nonrandom function that encodes the conditional specification of  $Z(\cdot)$ . We use the symbol  $\theta$  for this function and for the vector of covariance parameters since they both control the covariance function of  $Z(\cdot)$ . If for every  $\mathbf{x}$ ,  $\theta(\mathbf{x}, \mathbf{y})$  is zero for all but a finite number of locations  $\mathbf{y}$ , then the random field is said to be Markov. GMRFs can also be viewed as approximations to more general classes of Gaussian random fields (Rue and Tjelmeland, 2002). GMRFs are also used as a stage in other spatial hierarchical multiresolution models (Nychka et al., 2014) and non-Gaussian models (Rue et al., 2009). The most common approaches to achieving efficient Bayesian inference with latent GMRFs involve Markov chain Monte Carlo (MCMC) (Knorr-Held and Rue, 2002) or an integrated nested Laplace approximation (INLA) (Rue et al., 2009) to make inferences about  $\theta$  and  $\sigma^2$ . For large datasets, the usual MCMC methods involve repeated proposals of  $\mathbf{Z}$  to integrate out the latent GMRF.

INLA uses a clever rewriting of the likelihood for the GMRF parameters given  $\mathbf{Y}$  to avoid the need to integrate over the distribution of  $\mathbf{Z}$ . In principle, this same

approach could be used in MCMC as well. However, when using a finite sample of observations to make inferences about parameters controlling processes defined on infinite lattices, edge effects are known to introduce nonnegligible biases (Guyon, 1982). For GMRFs, the issue is that conditional distributions of observations on the boundary of the study region given only the observations on the interior cannot be easily expressed in terms of  $\theta$  and depend on the particular configuration of the observations. This is also true for observations on the interior whose neighbors are missing. A common approach to mitigating edge effects is to expand the lattice and treat observations on the expanded lattice as missing (Paciorek, 2007; Lindgren et al., 2011; Stroud et al., 2014; Guinness and Fuentes, 2014). This approach is generally approximate in nature, and the practitioner must decide how much lattice expansion is required to mitigate the edge effects. Stroud et al. (2014) describe an exact method, but Guinness and Fuentes (2014) argue that the tradeoff for exactness is a large number of missing values that must be imputed, slowing the convergence of iterative algorithms. Besag and Kooperberg (1995) give an approximate method for mitigating edge effects based on modifying an approximation to the precision matrix when the data form a complete rectangular grid. Dutta and Mondal (2014) use approximate likelihoods based on the h-likelihood (Henderson et al., 1959) but do not address edge effects.

We introduce methods for efficiently computing of the Gaussian likelihood for  $\theta$  and  $\sigma^2$  from  $\mathbf{Y}$  when the random field stationary, so when our methods can be applied, there is no need to impute missing values or expand the lattice to a larger domain. We do not impose any unrealistic boundary conditions on the spatial model—the likelihood is for the infinite lattice model observed on a finite lattice. The methods rely on the availability of the covariances, which we show can be computed efficiently and to very high precision with fast Fourier transform (FFT) algorithms, and we provide a theorem and a numerical study as support. The computational methods make use of circulant embedding techniques as well as sparse matrix algorithms. The computations are efficient in both memory allocation and floating point operations, as long as the observations form a nearly complete grid. The boundary of the grid need not be rectangular or even convex, and thus the methods can be used in diverse applications, particularly those in which the boundary of the study region is determined by irregular geography. Indeed, we apply our methods to a set of aerosol optical thickness values observed over the Red Sea, which has an irregular coastline and several islands over which the data are not reported. In Section 2, we describe our new methods for computing the exact likelihood, as well as discuss some existing likelihood approximations. Section 3 contains simulation and timing results showing cases in which the existing likelihood approximations fail due to boundary effects, and demonstrates the computational efficiency of our new methods. The analysis of the Red Sea data is in Section 4, where we highlight the role of the exact likelihood for both parameter estimation and model comparison, and where stochastic partial differential equation approximations are outperformed by much simpler independent blocks likelihood approximations. We conclude with a discussion in Section 5.

## 2 Efficient Gaussian Likelihood Calculations

In this section, we describe our theoretical results and methods for computing likelihoods for stationary GMRF models. For any  $\mathbf{x} \in \mathbb{Z}^2$ , define  $S_\theta(\mathbf{x})$  to be the set

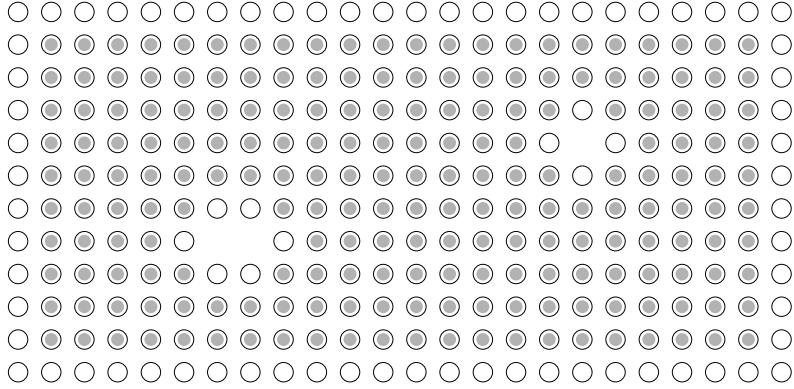


Figure 1: Illustration of the fully and partially neighbored observations with neighborhood set consisting of four adjacent grid points. Black circles indicate observation locations, and those with a gray interior are the fully neighbored locations.

$\{\mathbf{y} \in \mathbb{Z}^2 : \theta(\mathbf{x}, \mathbf{y}) \neq 0\}$ , which we refer to as the neighborhood set of  $\mathbf{x}$  under  $\theta$ . If the random field is Markov, then it is possible that all neighbors of an observation location are contained in the finite set of all observation locations.

**Definition:** Let  $1 \leq i \leq n$ . If  $S_\theta(\mathbf{x}_i) \subset \mathbf{x}_{1:n}$ , we say that  $\mathbf{x}_i$  is a *fully neighbored* location, and  $Z(\mathbf{x}_i)$  is a *fully neighbored* observation. If a location or observation is not fully neighbored, we say that it is *partially neighbored*.

We define  $m_n \leq n$  to be the number of partially neighbored observations among  $\mathbf{x}_{1:n}$ . The partially neighbored observations consist of those along the boundary of the study region and any observations on the interior of the study region that have at least one missing neighbor. Figure 1 contains an illustration. For example, if  $\mathbf{x}_{1:n}$  is a complete square grid of dimension  $(1000, 1000)$ , and each location's neighborhood set is the four adjacent grid points,  $m_n = 3996$ , while  $n = 10^6$ . Practical application of our methods depends on the following lemma, so we state it here first to give context to the theoretical results that follow.

**Lemma 1.** Let  $\mathbf{x}_{1:n}$  be a finite set of observation locations for the infinite GMRF  $Z(\cdot)$  specified by  $\theta$ , let  $Q_\theta = \Sigma_\theta^{-1}$ , the inverse of the covariance matrix for  $\mathbf{Z}$ , and let  $\mathbf{x}_i$  and  $\mathbf{x}_j$  be two observation locations. If either  $\mathbf{x}_i$  or  $\mathbf{x}_j$  is fully neighbored, then  $Q_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$ .

Let  $P$  be a permutation matrix for which  $P\mathbf{Z} = (\mathbf{Z}'_1, \mathbf{Z}'_2)'$  reorders the components of  $\mathbf{Z}$  so that  $\mathbf{Z}_1$  contains the  $m_n$  partially neighbored observations and  $\mathbf{Z}_2$  contains the  $n - m_n$  fully neighbored observations. Then define the block matrices

$$P\Sigma_\theta P' = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}, \quad PQ_\theta P' = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix}, \quad (2)$$

so that  $\Sigma_{11}$  the covariance matrix for  $\mathbf{Z}_1$ , and  $\Sigma_{22}$  is the covariance matrix for  $\mathbf{Z}_2$ . Since  $Q_{12}$ ,  $Q_{21}$  and  $Q_{22}$  all contain either a row or a column corresponding to a fully neighbored observation, they are all sparse and have entries given by Lemma 1. On the

other hand  $Q_{11}$  is not guaranteed to be sparse, and its entries cannot be exactly determined by  $\theta$ . Indeed, this issue of edge effects has been a topic of continued discussion throughout the Markov random field literature, see for example Besag and Kooperberg (1995), and Besag and Higdon (1999).

When  $\theta(\mathbf{x}, \mathbf{y})$  can be expressed as  $\eta(\mathbf{x} - \mathbf{y})$  for every  $\mathbf{x}$  and  $\mathbf{y}$ , the random field is said to be stationary. In this case, the covariance function has a spectral representation

$$K_\theta(\mathbf{x}, \mathbf{y}) = \int_{[0, 2\pi]^2} \frac{\exp(i\omega'(\mathbf{x} - \mathbf{y}))}{\sum_{\mathbf{h} \in \mathbb{Z}^2} \eta(\mathbf{h}) \exp(i\omega'\mathbf{h})} d\omega, \quad (3)$$

where the spectral density  $f_\theta(\omega) = (\sum_{\mathbf{h} \in \mathbb{Z}^2} \eta(\mathbf{h}) \exp(i\omega'\mathbf{h}))^{-1}$  is required to be non-negative for all  $\omega \in [0, 2\pi]^2$ . Assume that all of the observation locations fall inside a rectangular grid of size  $\mathbf{n} = (n_1, n_2)$ , and the axes have been defined so that  $n_1 < n_2$ . To compute the covariances, we follow a suggestion by Besag and Kooperberg (1995) to numerically evaluate the integrals in (3). We use the sums

$$K_\theta(\mathbf{x}, \mathbf{y}; J, \mathbf{n}) = \frac{4\pi^2}{n_1 n_2 J^2} \sum_{\mathbf{j} \in F_J} f_\theta(\omega_j) \exp(i\omega'_j(\mathbf{x} - \mathbf{y})), \quad (4)$$

where  $\omega_j = 2\pi/J(j_1/n_1, j_2/n_2)$  are Fourier frequencies on a grid  $F_J$  of size  $(n_1 J, n_2 J)$ . The following theorem gives a bound on the convergence rate of this method for calculating the covariances.

**Theorem 1.** *If  $f(\omega)$  is the spectral density for a Markov random field on  $\mathbb{Z}^2$  and is bounded above, then for  $J \geq 2$  and any  $p \in \mathbb{N}$ , there exists a constant  $C_p < \infty$  such that*

$$|K_\theta(\mathbf{x}, \mathbf{y}) - K_\theta(\mathbf{x}, \mathbf{y}; J, \mathbf{n})| \leq \frac{C_p}{(n_1 J)^p}.$$

We provide a proof in Appendix B. In other words, the error decays faster than a polynomial in  $(n_1 J)$  of degree  $p$  for every  $p \geq 1$ , so the sums converge quickly with  $J$ , and the error is smaller for a given  $J$  with larger  $n_1$ . For large observation grids and common GMRF models, the covariances can be computed to an absolute error of less than  $10^{-15}$  with  $J$  as small as 2 or 3. We provide a numerical study in Appendix A. FFT algorithms allow us to compute covariances at all of the required lags  $\mathbf{x} - \mathbf{y}$  in  $O((n_1 n_2 J^2) \log(n_1 n_2 J^2))$  floating point operations and  $O(n_1 n_2 J^2)$  memory.

The following proposition gives convenient expressions for the determinant and inverse of  $\Sigma_\theta$  that we use to efficiently evaluate the Gaussian loglikelihood in the case when the model does not contain a nugget.

**Proposition 1.** *If  $\Sigma_\theta$  and  $Q_\theta$  are partitioned as in (2), then*

$$(a) \quad \det(\Sigma_\theta) = \det(\Sigma_{11}) / \det(Q_{22})$$

$$(b) \quad \Sigma_\theta^{-1} = \begin{bmatrix} I & -\Sigma_{11}^{-1}\Sigma_{12} \\ 0 & I \end{bmatrix} \begin{bmatrix} \Sigma_{11}^{-1} & 0 \\ 0 & Q_{22} \end{bmatrix} \begin{bmatrix} I & 0 \\ -\Sigma_{21}\Sigma_{11}^{-1} & I \end{bmatrix}$$

Proposition 1 can be viewed as a matrix version of the decomposition of the likelihood  $p(\mathbf{z}_1, \mathbf{z}_2) = p(\mathbf{z}_1)p(\mathbf{z}_2|\mathbf{z}_1)$ , since  $\mathbf{Z}_1$  is normal with mean 0 and covariance matrix  $\Sigma_{11}$ , and  $\mathbf{Z}_2|\mathbf{Z}_1$  is normal with mean  $\Sigma_{21}\Sigma_{11}^{-1}\mathbf{Z}_1$  and covariance matrix  $Q_{22}^{-1}$ , which is the

Schur complement of  $\Sigma_{11}$ . The entries of  $\Sigma_{11}$  are obtained with an FFT, as in Theorem 1, and the entries of  $Q_{22}$  are known analytically. The various quantities required for evaluating the determinant and for multiplying  $\Sigma_\theta^{-1}$  by  $\mathbf{Y}$  can be easily computed after the Cholesky factorizations of  $\Sigma_{11}$  and  $Q_{22}$  have been obtained. Neither factorization is computationally demanding because  $\Sigma_{11}$ , though dense, is only of size  $m_n \times m_n$ , and  $Q_{22}$ , though large, is sparse and can be factored with sparse Cholesky algorithms. Multiplication of  $\Sigma_{21}$  or  $\Sigma_{12}$  by a vector is completed with circulant embedding techniques (Wood and Chan, 1994).

Including a nugget in the model increases the computational demands because  $(\Sigma_\theta + \sigma^2 I)^{-1}$  is typically dense, even if  $\Sigma_\theta^{-1}$  is sparse. Recent work on solving diagonally perturbed linear systems include Bellavia et al. (2011); Du et al. (2015); Erway and Marcia (2012); Soodhalter et al. (2014); Vecharynski and Knyazev (2013). Here, we exploit the fact that most of the inverse of the unperturbed matrix is known and sparse, and we provide methods for evaluating the determinant in addition to solving linear systems. The covariance matrix for  $\mathbf{Y}$  can be decomposed as

$$\Sigma_\theta + \sigma^2 I = \Sigma_\theta(I + \sigma^2 Q_\theta) = \Sigma_\theta(\sigma^{-2} I + Q_\theta)\sigma^2 I, \quad (5)$$

and thus the following proposition holds.

**Proposition 2.**

- (a)  $\det(\Sigma_\theta + \sigma^2 I) = \det(\Sigma_\theta) \det(I + \sigma^2 Q_\theta)$
- (b)  $(\Sigma_\theta + \sigma^2 I)^{-1} = (I + \sigma^2 Q_\theta)^{-1} Q_\theta$ .

Equation (5) and Proposition 2 can be viewed as a matrix version of the decomposition of the likelihood  $p(\mathbf{y}) = p(\mathbf{z})p(\mathbf{y}|\mathbf{z})/p(\mathbf{z}|\mathbf{y})$ , since the covariance matrix for  $\mathbf{Z}$  is  $\Sigma_\theta$ , the covariance matrix for  $\mathbf{Y}|\mathbf{Z}$  is  $\sigma^2 I$ , and the precision matrix for  $\mathbf{Z}|\mathbf{Y}$  is  $\sigma^{-2} I + Q_\theta$  (Lindgren et al., 2011). Exploiting this factorization requires one to know the dense matrix  $Q_{11}$ , but as a consequence of Theorem 1,  $Q_{11}$  can be computed as

$$Q_{11} = \Sigma_{11}^{-1} + Q_{12}Q_{22}^{-1}Q_{21},$$

which requires inversion of  $\Sigma_{11}$ ,  $m_n$  solves with the large sparse matrix  $Q_{22}$ , followed by a multiplication with  $Q_{12}$ , which is also sparse. In practice, we have found that the  $m_n$  solves are the most time-consuming step, but the solves can be completed in parallel.

In the examples we consider, it is possible to compute and store the Cholesky factors of  $Q_\theta$  and  $I + \sigma^2 Q_\theta$  if their rows and columns are appropriately reordered, so Proposition 2 can be used directly to compute the Gaussian likelihood in these cases. However, storing the Cholesky factors could be a computationally limiting constraint in some cases. The following proposition gives expressions for the determinant and inverse of  $\Sigma_\theta + \sigma^2 I$  that can be used to evaluate the Gaussian loglikelihood function without storing any dense matrices larger than size  $m_n \times m_n$ .

**Proposition 3.** *Let  $A := I + \sigma^2 Q_\theta$  and  $B := A^{-1}$  be partitioned as in (2). Then*

- (a)  $\det(\Sigma_\theta + \sigma^2 I) = \det(\Sigma_\theta) \det(A_{22}) \det(B_{11}^{-1})$
- (b)  $(\Sigma_\theta + \sigma^2 I)^{-1} = \begin{bmatrix} I & -\sigma^2 Q_{12}A_{22}^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} B_{11} & 0 \\ 0 & A_{22}^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -\sigma^2 A_{22}^{-1}Q_{21} & I \end{bmatrix} Q_\theta,$

and  $B_{11}^{-1} = I + \sigma^2 Q_{11} - \sigma^4 Q_{12}A_{22}^{-1}Q_{21}$ .

Since  $B_{11}$  is the same size as  $\Sigma_{11}$  and  $A_{22}$  is as sparse as  $Q_{22}$ , most of the calculations are analogous to those that appear in Proposition 1, so we omit most of the details here, noting however that multiplications with  $Q_\theta$  are fast since it is mostly sparse, and we first form the matrix  $B_{11}^{-1}$ , which is the Schur complement of  $A_{22}$ , then factor it to compute its determinant and solve linear systems with it. Further, when forming  $B_{11}^{-1}$ , we do not need to store the entire matrix  $A_{22}^{-1}Q_{21}$ , which is dense and of size  $n \times m_n$ . Rather, we complete the solves involving each column  $Q_{21}[:, j]$  individually, either in sequence or in parallel, only storing  $Q_{12}A_{22}^{-1}Q_{21}[:, j]$ .

## 2.1 Kriging and Conditional Simulations

The inverse results in Propositions 1, 2, and 3 can also be used to perform Kriging of the data to unobserved locations, since the computationally limiting step in Kriging is solving a linear system involving the covariance matrix and the data. Let  $\mathbf{Y}_0$  be a vector of unobserved values that we wish to predict using the vector of observations  $\mathbf{Y}$ . The joint distribution of the two vectors is

$$\begin{bmatrix} \mathbf{Y} \\ \mathbf{Y}_0 \end{bmatrix} \sim N\left(\begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu}_0 \end{bmatrix}, \begin{bmatrix} \Sigma_\theta + \sigma^2 I & \Sigma_0 \\ \Sigma_0^T & \Sigma_{00} + \sigma^2 I \end{bmatrix}\right),$$

where  $\boldsymbol{\mu}$  and  $\boldsymbol{\mu}_0$  are the two mean vectors,  $\Sigma_0$  is the cross covariance matrix between  $\mathbf{Y}$  and  $\mathbf{Y}_0$ , and  $\Sigma_{00} + \sigma^2 I$  is the covariance matrix for the unobserved values. Then the conditional expectation of  $\mathbf{Y}_0$  given  $\mathbf{Y}$  is

$$E(\mathbf{Y}_0 | \mathbf{Y}) = \boldsymbol{\mu}_0 + \Sigma_0^T (\Sigma_\theta + \sigma^2 I)^{-1} (\mathbf{Y} - \boldsymbol{\mu})$$

The inverse results are used to solve the linear system with  $\Sigma_\theta + \sigma^2 I$ , and the forward multiplication with  $\Sigma_0$  can be computed with circulant embedding techniques.

Conditional simulations of  $\mathbf{Y}_0$  given  $\mathbf{Y}$  do not require much additional computational effort. The conditional simulations consist of the conditional expectation (computed above), added to a simulation of a conditional residual with covariance matrix  $\Sigma_{00} + \sigma^2 I - \Sigma_0^T (\Sigma_\theta + \sigma^2 I)^{-1} \Sigma_0$ . The conditional residual can be formed with standard methods that are not more computationally demanding than the original Kriging computations, as long as one can generate unconditional simulations. See for example Chiles and Delfiner (2012, Section 7.3.1). We use circulant embedding for the unconditional simulations.

## 2.2 Existing Approximations to $Q_\theta$

We consider some existing likelihood approximations as competitors to our methods for computing the exact likelihood. The likelihood approximations replace  $Q_\theta = \Sigma_\theta^{-1}$  with an approximation  $\tilde{Q}_\theta$ , with the various approximations differing in how  $\tilde{Q}_\theta[i, j]$  is defined when both  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are partially neighbored, that is, how  $\tilde{Q}_{11}$  is defined. Thus each approximation can be viewed as a method for dealing with edge effects and missing interior values. We may prefer to use a likelihood approximation if it is faster to compute and provides similar parameter estimates to the exact maximum likelihood parameters. We describe three approximations here and evaluate their effectiveness in Section 3.

*No Adjustment:* The simplest approximation defines  $\tilde{Q}_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$  for every  $i$  and  $j$ . Rue and Held (2005) prove that  $\tilde{Q}_\theta$  is positive definite when defined in this way.

*Precision Adjustment:* When the conditional specification satisfies the diagonal dominance criterion

$$\theta(\mathbf{x}_i, \mathbf{x}_i) = \lambda \sum_{\mathbf{x}_j \in \mathbb{Z}^2 \setminus \mathbf{x}_i} |\theta(\mathbf{x}_i, \mathbf{x}_j)|$$

with  $\lambda \in [0, 1]$ , we set  $\tilde{Q}_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$  for  $i \neq j$  and

$$\tilde{Q}_\theta[i, i] = \lambda \sum_{\mathbf{x}_j \in \mathbf{x}_{1:n} \setminus \mathbf{x}_i} |\theta(\mathbf{x}_i, \mathbf{x}_j)|$$

when both  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are partially neighbored. Then  $\tilde{Q}_\theta$  is symmetric and diagonally dominant, and thus positive definite.

*Periodic (Toroidal) Adjustment:* Let  $\mathbf{x}_{1:n}$  be a complete rectangular subset of the integer lattice of dimensions  $\mathbf{n} = (n_1, n_2)$ , so that  $n = n_1 n_2$ . Then we set

$$\begin{aligned} \tilde{Q}_\theta[i, j] &= \theta(\mathbf{x}_i - \mathbf{x}_j + \mathbf{k} \circ \mathbf{n}), \\ k_\ell &= \arg \min_{k \in \{-1, 0, 1\}} |x_{\ell i} - x_{\ell j} + kn_\ell| \end{aligned}$$

for  $\ell = 1, 2$ , where  $\mathbf{k} = (k_1, k_2)$ , and  $\mathbf{k} \circ \mathbf{n} = (k_1 n_1, k_2 n_2)$ .

### 3 Simulation and Timing Experiments

In this section we study the maximum likelihood and maximum approximate likelihood parameter estimates in simulation and timing experiments. For the simulations, we choose to consider the case of no nugget, so that measurement error does not obscure the importance of edge effects. This allows us to study how the approximate edge correction techniques perform compared to inference using the exact likelihood. Existing MCMC and INLA methods are designed for Bayesian inference, and thus are not applicable here. We do, however, consider the performance of INLA in the data analysis in Section 4.

We consider spectral densities of the form

$$f_\theta(\boldsymbol{\omega}) = \tau^{-2} (\kappa^2 + 4 - e^{i\omega_1} - e^{-i\omega_1} - e^{i\omega_2} - e^{-i\omega_2})^{-\nu-1} \quad (6)$$

where  $\theta = (\tau, \kappa, \nu)$  with  $\tau, \kappa > 0$  and  $\nu = 0, 1, 2, \dots$ . Lindgren et al. (2011) showed that the spectral density in (6) is explicitly linked to the spectral density of the Matérn covariance function with integer smoothness parameter  $\nu$  and inverse range parameter  $\kappa$ . When  $\nu = 0$ , we have  $\theta(\mathbf{0}) = \tau^2(\kappa^2 + 4)$ ,  $\theta(\mathbf{h}) = -\tau^2$  when  $\mathbf{h} = (1, 0), (-1, 0), (0, 1)$ , or  $(0, -1)$ , and 0 otherwise. The  $\nu = 0$  model has been called a symmetric first order autoregression (Besag and Kooperberg, 1995), although an alternative parameterization is more common.

The following simulation studies demonstrate that edge effects adversely affect parameter estimation, especially when the spatial correlation is strong. We consider the cases of  $\nu = 0$  and  $\nu = 1$  with  $\kappa \in \{1/5, 1/10, 1/20\}$  and We conduct 100 simulations

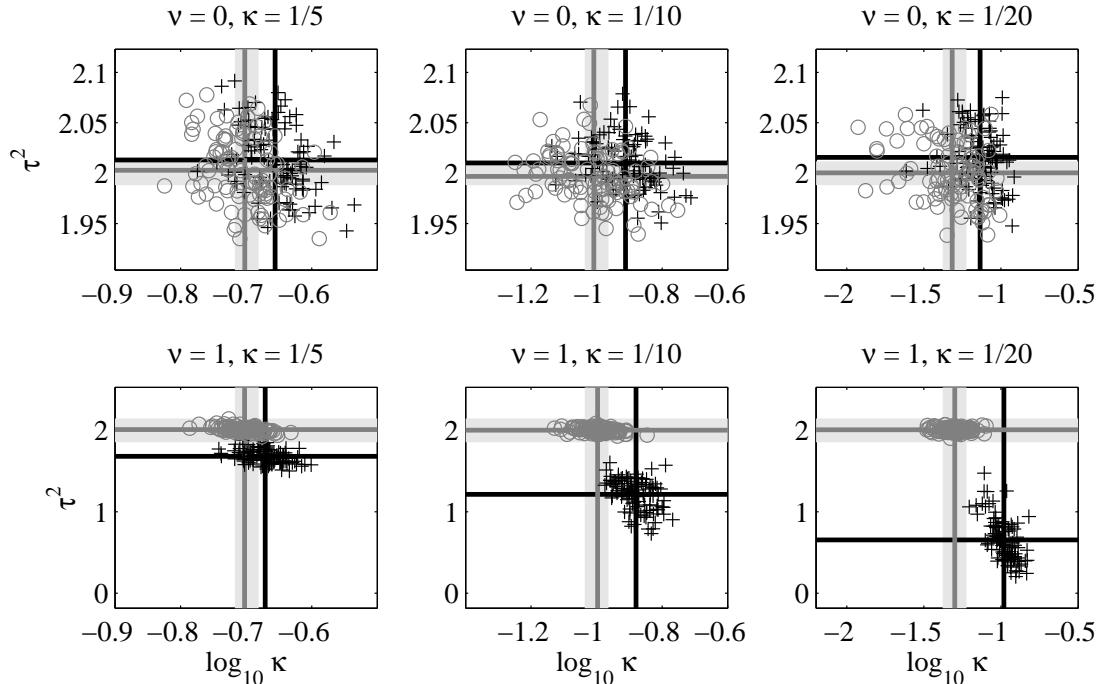


Figure 2: Simulation results for maximum approximate likelihood estimates (black crosses) and maximum exact likelihood estimates (dark gray circles) for 100 simulated fields at each parameter combination. The thin solid lines indicate the sample means of the parameter estimates, with the sample mean of  $\kappa$  taken on the log scale. The thick light gray lines indicate the true parameter values.

on a grid of size  $(100, 100)$  for each of the six parameter combinations. We compare the maximum approximate likelihood estimates to the maximum likelihood estimates, computed with our efficient methods. For the  $\nu = 0$  case, we find that the precision adjustment provides the best approximation. For the  $\nu = 1$  case, we find that the periodic adjustment provides the best approximation, although no adjustment performs similarly. The results are presented in Figure 2. In Appendix C, we provide a figure containing simulated fields with these six parameter combinations. The precision adjustment introduces a small bias when  $\nu = 0$ . When  $\nu = 1$ , the periodic adjustment performs poorly in every case, with the bias of the estimates increasing as  $\kappa$  decreases. The maximum likelihood estimates are nearly unbiased in every case.

When there is no nugget in the model, the exact likelihoods not only produce more reliable parameter estimates, they often do not impose a significant additional computational burden over the calculation of the approximate likelihoods. In Table 1, we present the results of a timing study for the  $\nu = 0$  and  $\nu = 1$  cases using no nugget and “no adjustment” neighborhood structure for the approximate likelihoods, which is the sparsest approximation. We compute the exact loglikelihood both with and without a small nugget. We use complete square lattices of increasing size, and time the computations using the “clock” and “etime” functions in Matlab. In both the  $\nu = 0$  and the  $\nu = 1$  case with no nugget, the approximate likelihood is only 1.25 times faster than the exact likelihood for the largest grid. Adding a nugget increases the

| $n$     | $\nu = 0$      |                |                |                   | $\nu = 1$      |                |                   |
|---------|----------------|----------------|----------------|-------------------|----------------|----------------|-------------------|
|         | Approximate    |                | Exact          | Exact             | Approximate    | Exact          | Exact             |
|         | $\sigma^2 = 0$ | $\sigma^2 = 0$ | $\sigma^2 = 0$ | $\sigma^2 = 0.01$ | $\sigma^2 = 0$ | $\sigma^2 = 0$ | $\sigma^2 = 0.01$ |
| $100^2$ | 0.05           | 0.06           | 0.16           |                   | 0.06           | 0.09           | 0.89              |
| $150^2$ | 0.09           | 0.13           | 0.55           |                   | 0.22           | 0.27           | 2.38              |
| $200^2$ | 0.22           | 0.33           | 1.20           |                   | 0.55           | 0.64           | 6.90              |
| $250^2$ | 0.47           | 0.56           | 2.17           |                   | 1.25           | 1.41           | 15.06             |
| $300^2$ | 0.80           | 0.95           | 3.72           |                   | 1.95           | 2.45           | 22.72             |

Table 1: Time in seconds for one evaluation of the approximate likelihood and the exact likelihood, with and without a nugget term.

time to compute the exact likelihood, but we note that the largest computations are infeasible without our new methods, due to memory constraints on naive methods, and we have not parallelized the solves required for constructing  $Q_{11}$ . All computations are completed using Matlab R2013a on an Intel Core i7-4770 CPU (3.40GHz) with 32 GB of RAM.

## 4 Application to Satellite Data

The NASA Aqua satellite carries a moderate resolution imaging spectrometer (MODIS) capable of acquiring radiance data at high spatial resolution. The data are processed in order to obtain derived measurements of physical quantities relevant to land, ocean, and atmospheric dynamics, at various timescales. We analyze a set of aerosol optical thickness (AOT) data from Autumn 2014 over the Red Sea. The data were downloaded from NASA’s OceanColor project website. Aerosol optical thickness is a unitless quantity describing the radiance attenuation caused by aerosols in the atmosphere. In Figure 3, we plot the data, which consist of 21,921 observations forming an incomplete regular grid. Even when gridded, satellite data often contain many missing values due to lack of coverage, atmospheric disturbances, or the quantity not being defined in certain regions (e.g. sea surface temperature over land).

Our methods for computing the likelihood are useful both for estimating parameters and for comparing models, allowing us to conduct model selection and to compare various methods for estimating parameters. We highlight both of these uses in the analysis of the AOT data. We model the data as

$$Y(\mathbf{x}) = \mu + Z(\mathbf{x}) + \varepsilon(\mathbf{x}),$$

where  $\mu \in \mathbb{R}$  is unknown and nonrandom,  $Z(\cdot)$  is a GMRF with spectral density as in (6), with  $\nu = 1$ , and unknown nonrandom parameters  $\kappa, \tau > 0$ , and  $\varepsilon(\mathbf{x}) \sim$  i.i.d.  $N(0, \sigma^2)$ . For this choice of  $\nu = 1$ , and this set of observation locations, we have  $m_n = 3,469$ . We estimate  $\mu, \tau, \kappa$ , and  $\sigma^2$  by maximum likelihood, being careful to profile out  $\mu$  and  $\tau$ , which have closed form expressions for the values that maximize the likelihood given fixed values of the other parameters. We also estimate  $\mu, \tau$ , and  $\kappa$ , fixing  $\sigma^2 = 0$ , allowing for a comparison between models with and without a nugget effect. Adding a nugget increased the loglikelihood by 594 units, indicating that the

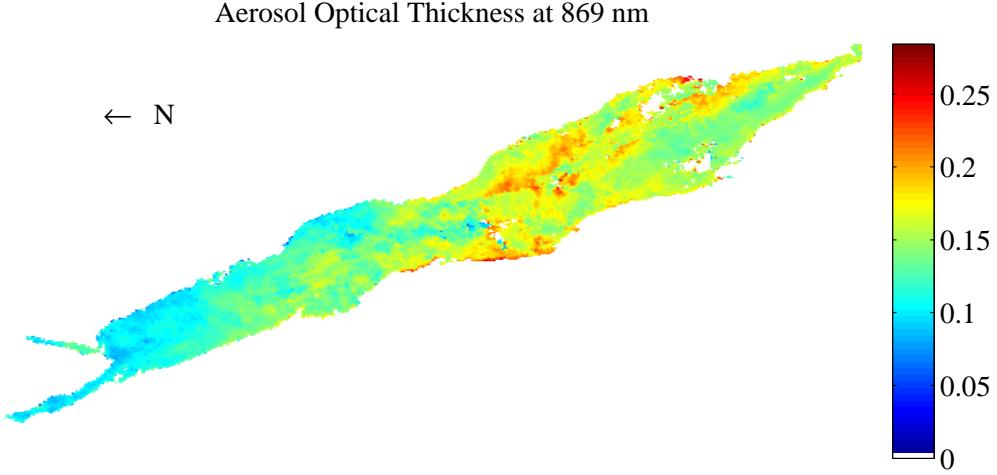


Figure 3: Aerosol optical thickness in the Red Sea in Autumn 2014. White pixels are missing values.

nugget significantly improves the model. The parameter estimates and loglikelihoods for these two models are given in Table 2.

Having access to the exact likelihood allows for likelihood ratio comparisons, a convenient way to compare parameter estimates found using various approximate methods. We estimate parameters with two well-known likelihood approximations: the independent blocks likelihood approximation (Stein, 2014), and an approximation based on an incomplete inverse Cholesky factorization (Vecchia, 1988; Stein et al., 2004). We also estimate parameters using the R-INLA software ([www.r-inla.org](http://www.r-inla.org)), described in Rue et al. (2009) and Martins et al. (2013), using a stochastic partial differential equation (SPDE) approximation (Lindgren et al., 2011). Since INLA is a Bayesian method, it does not return maximum likelihood estimates, but it is worthwhile to consider whether the INLA maximum posterior parameter estimates nearly maximize the likelihood, since INLA is designed specifically to estimate models with a latent GMRF.

There are several decisions that must be made when using the approximate methods. With the independent blocks method, we must choose the size and shape of the blocks. We partition the Red Sea according to the orientation in Figure 4, so that each block has roughly 1600 observations, allowing the covariance matrices for the individual blocks to be stored and factored using standard dense methods. With the incomplete Cholesky method, we choose smaller blocks, always using the adjacent block in the northwest direction as the conditioning set. In R-INLA, there are several decisions that must be made, including whether to force the nodes of the SPDE mesh to coincide with the observation locations, and how fine a mesh to use. And of course the priors must be chosen; we used the default priors. We did not enforce the mesh nodes to coincide with the grid. We tried values of 2, 4, and 6 for the “max.edge” option, with smaller values of max.edge giving finer meshes and better parameter estimates.

The results of the model fitting are summarized in Table 2. The use of the exact likelihood to estimate the model delivers estimates in just 8.40 minutes, which is a major step forward for exact likelihood methods for a dataset of this size. The time required for INLA estimates depends on the resolution of the mesh, with finer meshes

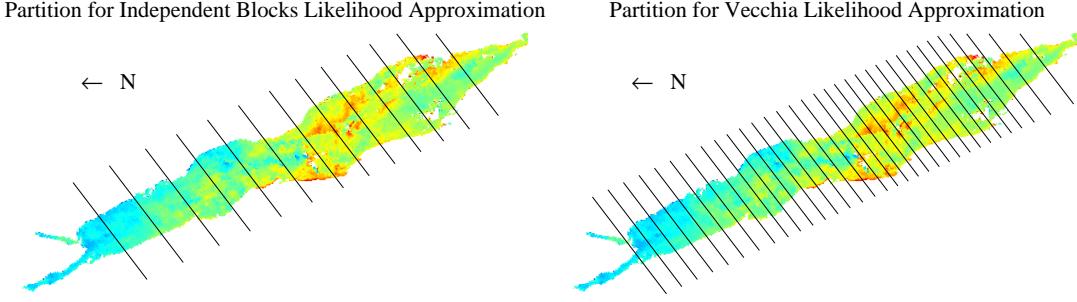


Figure 4: Partitions used for the independent blocks approximation and incomplete inverse Cholesky approximation (Vecchia Approximation).

| method             | $\mu$  | $\tau$ | $\kappa$ | $\sigma$ | $\Delta\text{loglik}$ | time (min) |
|--------------------|--------|--------|----------|----------|-----------------------|------------|
| exact with nugget  | 0.1464 | 51.25  | 0.2095   | 0.0048   | 0                     | 8.40       |
| inc. inv. Chol.    | 0.1468 | 49.29  | 0.2094   | 0.0050   | -33.4                 | 5.23       |
| ind. blocks        | 0.1485 | 49.12  | 0.2054   | 0.0051   | -54.3                 | 4.15       |
| INLA, max.edge = 2 | 0.1471 | 58.19  | 0.1715   | 0.0061   | -158.9                | 33.2       |
| exact no nugget    | 0.1462 | 33.63  | 0.3648   | 0        | -594.0                | 0.41       |
| INLA, max.edge = 4 | 0.1490 | 79.02  | 0.1126   | 0.0077   | -795.4                | 3.23       |
| INLA, max.edge = 6 | 0.1497 | 88.49  | 0.0855   | 0.0085   | -1344.1               | 1.23       |

Table 2: Results of analysis of Red Sea data. Loglikelihood differences are from maximum likelihood model.

requiring more time. Estimating the parameters with the finest mesh took 33.2 minutes, nearly four times longer than generating the maximum likelihood estimates. The coarser meshes were faster than maximum likelihood, but the quality of the parameter estimates found with INLA depends on the mesh, with coarser meshes producing worse parameter estimates. The finest mesh that ran faster than maximum likelihood gave estimates whose loglikelihood is 795 units below the maximum loglikelihood, and decreasing the resolution further gave a loglikelihood more than 1300 units below the maximum loglikelihood, so it appears that the decisions made in constructing the mesh are very important. Comparisons with INLA should be made with caution, since it is a Bayesian method, but it is interesting to note that, in terms of loglikelihoods, the finest mesh is outperformed by the much simpler independent blocks approximation, which required much less computational effort. The incomplete inverse Cholesky approximation was better still and slightly slower than independent blocks. In INLA, we tried experimenting with different priors and did not see much dependence on the priors.

## 5 Discussion

The Gaussian Markov random field, especially when paired with a nugget effect, is a powerful tool for modeling spatial data. This paper outlines how the likelihood for a GMRF plus a nugget can be computed when the GMRF model is stationary and the

data locations form a nearly complete regular grid. Our new methods naturally handle edge effects, which we demonstrate can cause serious problems for estimating parameters if not treated carefully. The availability of the likelihood allows for both frequentist and Bayesian inference, and it allows for likelihood ratio comparisons among parameter estimates obtained using various approximate methods. Our methods for computing the likelihood require forming the covariance matrix for the partially neighbored observations, and we show that the covariances can be computed efficiently and highly accurately with FFT algorithms. We make use of decompositions that exploit the fact that the inverse of the covariance matrix is mostly known and sparse. The likelihood calculations can be considered exact in the same sense that any calculation involving Matérn covariances is considered exact, since evaluation of the Matérn covariance function involves series expansions for the modified Bessel function.

The computational limits of our methods are determined by  $m_n$ , the number of partially neighbored observations, since we must construct and manipulate dense matrices of size  $m_n \times m_n$ . For complete square grids, this allows the computation of the likelihoods with more than one million observations. However, our methods are not able to handle situations in which  $m_n$  is much larger than  $10^4$ , which arises when the observation grid has many scattered missing values. In this situation, using an iterative estimation method that relies on conditionally simulating the missing values given the data (Stroud et al., 2014; Guinness and Fuentes, 2014) is appropriate. Then the methods described in this paper can be used to calculate the likelihood for the entire imputed dataset.

### Acknowledgements

The first author acknowledges support from the National Science Foundation's Division of Mathematical Sciences under grant number 1406016. The second author acknowledges the support from the XDATA Program of the Defense Advanced Research Projects Agency (DARPA), administered through Air Force Research Laboratory contract FA8750-12-C-0323 FA8750-12-C-0323.

## A Numerical Study

We demonstrate that the stationary covariances for the random fields that we consider can be computed to high accuracy with small values of  $J = 2$  or  $3$ . To study the calculations, for each value of  $J \in \{1, 2, 3, 4, 5\}$ , we compute the approximation in (4), and we record

$$\Delta_J = \max_{\mathbf{h}} \left| K_\theta(\mathbf{x}, \mathbf{x} + \mathbf{h}; J, \mathbf{n}) - K_\theta(\mathbf{x}, \mathbf{x} + \mathbf{h}; J + 1, \mathbf{n}) \right|,$$

where the maximum absolute difference is taken over  $\mathbf{h}$  on a grid of size  $\mathbf{n} = (100, 100)$ , for  $\nu \in \{0, 1\}$ , and  $\kappa \in \{1/5, 1/10, 1/20\}$ , in order to represent a wide set of models. The results are plotted in Figure 5. We see that the calculations converge very rapidly, with  $\Delta_3$  less than  $10^{-10}$  for every parameter combination.

## B Proofs

**Theorem 1.** *If  $f(\omega)$  is the spectral density for a Markov random field on  $\mathbb{Z}^2$  and is bounded above, then for  $J \geq 2$  and any  $p \in \mathbb{N}$ , there exists a constant  $C_p < \infty$  such*

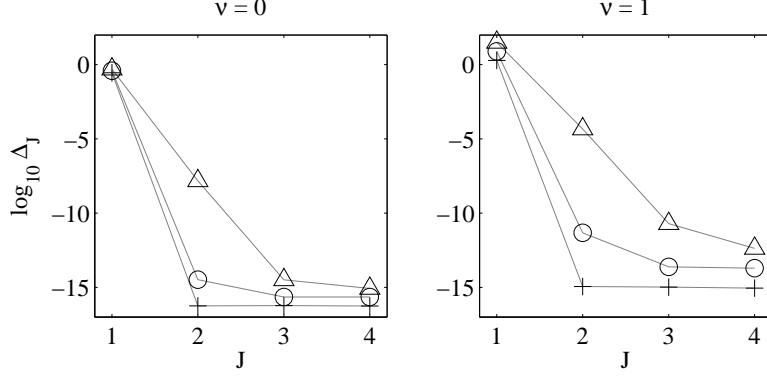


Figure 5: (+) represents  $\kappa = 1/5$ , (o) represents  $\kappa = 1/10$ , and ( $\Delta$ ) represents  $\kappa = 1/20$ .

that

$$|K_\theta(\mathbf{x}, \mathbf{y}) - K_\theta(\mathbf{x}, \mathbf{y}; J)| \leq \frac{C_p}{(n_1 J)^p}.$$

*Proof.* Since the random field is Markov,  $1/f(\omega)$  can be written as a finite sum of complex exponentials and must be infinitely differentiable, and thus  $f(\omega)$  is also infinitely differentiable, since it is assumed to be bounded above. Since the process is stationary, we write  $\mathbf{h} = \mathbf{x} - \mathbf{y}$ , and  $K(\mathbf{x}, \mathbf{y}) = K(\mathbf{h})$  to simplify the expressions. Guinness and Fuentes (2014) proved that

$$K(\mathbf{h}; J, \mathbf{n}) = \sum_{\mathbf{k} \in \mathbb{Z}^2} K(\mathbf{h} + (J\mathbf{n} \circ \mathbf{k})),$$

where  $\mathbf{n} \circ \mathbf{k} = (n_1 k_1, n_2 k_2)$ . Thus

$$K_\theta(\mathbf{h}) - K_\theta(\mathbf{h}; J, \mathbf{n}) = \sum_{\mathbf{k} \neq 0} K(\mathbf{h} + (J\mathbf{n} \circ \mathbf{k})).$$

Since  $f(\omega)$  is  $p \geq 3$  times continuously differentiable, there exists constant  $M_p < \infty$  such that  $|K(\mathbf{r})| \leq M_p \|\mathbf{r}\|^{-p}$ , where  $\|\cdot\|$  is Euclidean distance (Körner, 1989, Lemma 9.5). Then the error can be bounded by

$$\begin{aligned} |K_\theta(\mathbf{h}) - K_\theta(\mathbf{h}; J, \mathbf{n})| &\leq \sum_{\mathbf{k} \neq 0} |K(\mathbf{h} + (J\mathbf{n} \circ \mathbf{k}))| \\ &\leq 3M_p ((J-1)n_1)^{-p} + \sum_{j=1}^{\infty} 4(2j+1)M_p(jJn_1)^{-p} \\ &= 3M_p ((J-1)n_1)^{-p} + 4M_p(Jn_1)^{-p} \sum_{j=1}^{\infty} (2j+1)(j)^{-p} \\ &\leq 3M_p ((J/2)n_1)^{-p} + 4M_p(Jn_1)^{-p} \sum_{j=1}^{\infty} (2j+1)(j)^{-p} \\ &\leq \left[ 3M_p (1/2)^{-p} + 4M_p \sum_{j=1}^{\infty} (2j+1)(j)^{-p} \right] (Jn_1)^{-p} \end{aligned}$$

Setting  $C_p$  equal to the quantity in square brackets establishes the bound for  $p \geq 3$ . The  $p = 1$  and  $p = 2$  cases follow by setting  $C_1 = C_3$  and  $C_2 = C_3$ , since  $(Jn_1)^{-3} < (Jn_1)^{-2} < (Jn_1)^{-1}$   $\square$

We believe the following are known results but were unable to find proofs, so we prove them here for completeness. Lemma 2 is used in the proof of Lemma 1.

**Lemma 1.** *Let  $\mathbf{x}_{1:n}$  be a finite set of observation locations for the infinite GMRF  $Z(\cdot)$  specified by  $\theta$ , let  $Q_\theta = \Sigma_\theta^{-1}$ , the inverse of the covariance matrix for  $\mathbf{Z}$ , and let  $\mathbf{x}_i$  and  $\mathbf{x}_j$  be two observation locations. If either  $\mathbf{x}_i$  or  $\mathbf{x}_j$  is fully neighbored, then  $Q_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$ .*

*Proof.* Using the form of the density function of the multivariate normal, the density function of  $Z(\mathbf{x}_i)$  given  $\mathbf{Z}_{-i}$  ( $\mathbf{Z}$  with the  $i$ th element removed) is

$$\begin{aligned} p(z(\mathbf{x}_i)|\mathbf{z}_{-i}) &\propto \exp\left(-\frac{1}{2}\mathbf{z}'Q_\theta\mathbf{z}\right) \\ &\propto \exp\left(-\frac{1}{2}z(\mathbf{x}_i)^2Q_\theta[i, i] - z(\mathbf{x}_i)\sum_{j \neq i}Q_\theta[i, j]z(\mathbf{x}_j)\right), \end{aligned} \quad (7)$$

where  $\mathbf{z} = (z(\mathbf{x}_1), \dots, z(\mathbf{x}_n))$  is the argument for the joint density for  $\mathbf{Z}$ , and  $\mathbf{z}_{-i}$  is  $\mathbf{z}$  with the  $i$ th element removed. Suppose, without loss of generality, that  $\mathbf{x}_i$  is fully neighbored. Then, using the conditional specification in (1) and the form of the univariate normal distribution, the density function of  $Z(\mathbf{x}_i)$  given  $\mathbf{Z}_{-i}$  is

$$p(z(\mathbf{x}_i)|\mathbf{z}_{-i}) \propto \exp\left(-\frac{1}{2}z(\mathbf{x}_i)^2\theta(\mathbf{x}_i, \mathbf{x}_i) - z(\mathbf{x}_i)\sum_{j \neq i}\theta(\mathbf{x}_i, \mathbf{x}_j)z(\mathbf{x}_j)\right). \quad (8)$$

Comparing (7) and (8) proves that  $Q_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$ , since (7) and (8) must be equal for any  $\mathbf{z}_{-i} \in \mathbb{R}^{n-1}$ . The symmetry of  $Q_\theta$  and  $\theta$  (Lemma 2) implies that  $Q_\theta[j, i] = Q_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j) = \theta(\mathbf{x}_j, \mathbf{x}_i)$ , which establishes the stated result.  $\square$

**Lemma 2.** *If  $\theta$  specifies a valid GMRF,  $\theta(\mathbf{x}_i, \mathbf{x}_j) = \theta(\mathbf{x}_j, \mathbf{x}_i)$  for all  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{Z}^2$ .*

*Proof.* Let  $\mathbf{x}_{1:n}$  be a set of observation locations that includes  $\mathbf{x}_i$ ,  $\mathbf{x}_j$ ,  $S_\theta(\mathbf{x}_i)$ , and  $S_\theta(\mathbf{x}_j)$ . The conditional density of  $Z(\mathbf{x}_i)$  given  $\mathbf{Z}_{-i}$  is

$$\begin{aligned} p(z(\mathbf{x}_i)|\mathbf{z}_{-i}) &\propto \exp\left(-\frac{1}{2}\mathbf{z}'Q_\theta\mathbf{z}\right) \\ &\propto \exp\left(-\frac{1}{2}z(\mathbf{x}_i)^2Q_\theta[i, i] - z(\mathbf{x}_i)\sum_{j \neq i}Q_\theta[i, j]z(\mathbf{x}_j)\right). \end{aligned} \quad (9)$$

Using the form of the univariate normal distribution and the conditional mean and variance of  $Z(\mathbf{x}_i)$  given its neighbors, the conditional density of  $Z(\mathbf{x}_i)$  given  $\mathbf{x}_{1:n}$  is

$$p(z(\mathbf{x}_i)|\mathbf{z}_{-i}) \propto \exp\left(-\frac{1}{2}z(\mathbf{x}_i)^2\theta(\mathbf{x}_i, \mathbf{x}_i) - z(\mathbf{x}_i)\sum_{j \neq i}\theta(\mathbf{x}_i, \mathbf{x}_j)z(\mathbf{x}_j)\right). \quad (10)$$

Comparing (9) and (10) and equating coefficients gives  $Q_\theta[i, j] = \theta(\mathbf{x}_i, \mathbf{x}_j)$ . Repeating the same steps for  $\mathbf{x}_j$  gives  $Q_\theta[j, i] = \theta(\mathbf{x}_j, \mathbf{x}_i)$ . The symmetry of  $Q_\theta$  implies that  $\theta(\mathbf{x}_i, \mathbf{x}_j) = \theta(\mathbf{x}_j, \mathbf{x}_i)$ .  $\square$

## C Example Plots of Simulated Data

Figure 6 shows realizations from the six models considered in the simulation study from Section 3.

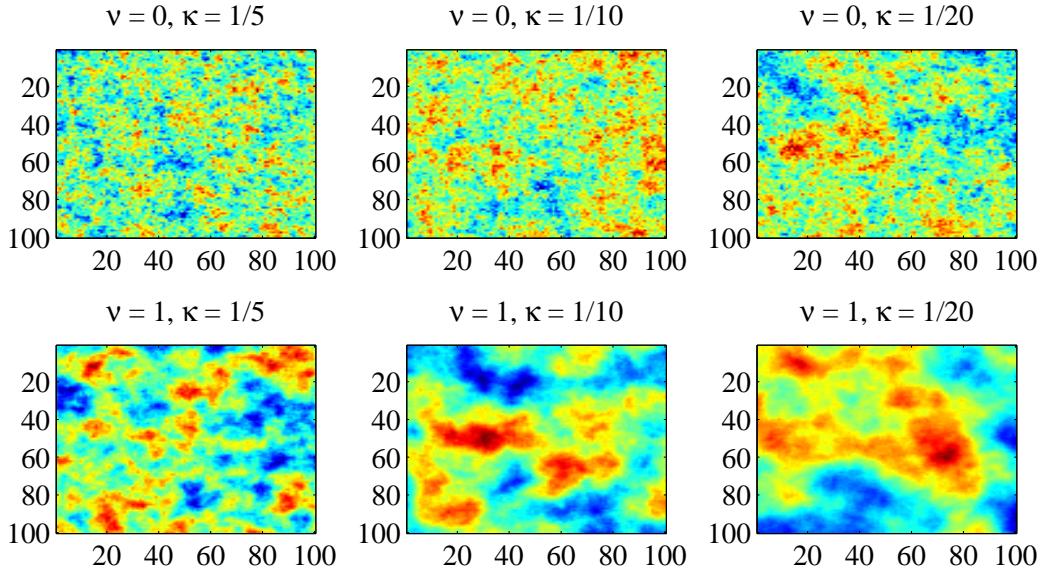


Figure 6: Examples of simulated data for the six parameter combinations.

## References

- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical modeling and analysis for spatial data*. Crc Press.
- Bellavia, S., De Simone, V., di Serafino, D., and Morini, B. (2011). Efficient preconditioner updates for shifted linear systems. *SIAM J. Sci. Comput.*, 33(4):1785–1809.
- Besag, J. and Higdon, D. (1999). Bayesian analysis of agricultural field experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(4):691–746.
- Besag, J. and Kooperberg, C. (1995). On conditional and intrinsic autoregressions. *Biometrika*, 82(4):733–746.
- Chiles, J.-P. and Delfiner, P. (2012). *Geostatistics: modeling spatial uncertainty*. John Wiley & Sons, 2 edition.
- Du, L., Sogabe, T., and Zhang, S.-L. (2015). IDR( $s$ ) for solving shifted nonsymmetric linear systems. *J. Comput. Appl. Math.*, 274:35–43.
- Dutta, S. and Mondal, D. (2014). An h-likelihood method for spatial mixed linear models based on intrinsic auto-regressions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.

- Erway, J. B. and Marcia, R. F. (2012). Limited-memory BFGS systems with diagonal updates. *Linear Algebra Appl.*, 437(1):333–344.
- Guinness, J. and Fuentes, M. (2014). Circulant embedding of approximate covariances for inference from Gaussian data on large lattices. *NC State Department of Statistics Technical Reports*, (2657).
- Guyon, X. (1982). Parameter estimation for a stationary process on a d-dimensional lattice. *Biometrika*, 69(1):95–105.
- Henderson, C. R., Kempthorne, O., Searle, S. R., and Von Krosigk, C. (1959). The estimation of environmental and genetic trends from records subject to culling. *Biometrics*, 15(2):192–218.
- Knorr-Held, L. and Rue, H. (2002). On block updating in Markov random field models for disease mapping. *Scandinavian Journal of Statistics*, 29(4):597–614.
- Körner, T. W. (1989). *Fourier analysis*. Cambridge University Press.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.
- Martins, T. G., Simpson, D., Lindgren, F., and Rue, H. (2013). Bayesian computing with INLA: new features. *Computational Statistics & Data Analysis*, 67:68–83.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2014). A multi-resolution Gaussian process model for the analysis of large spatial data sets. *Journal of Computational and Graphical Statistics*, (just-accepted):00–00.
- Paciorek, C. J. (2007). Bayesian smoothing with Gaussian processes using Fourier basis functions in the spectralgp package. *Journal of Statistical Software*, 19(2):nihpa22751.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC Press.
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):319–392.
- Rue, H. and Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields. *Scandinavian Journal of Statistics*, 29(1):31–49.
- Soodhalter, K. M., Szyld, D. B., and Xue, F. (2014). Krylov subspace recycling for sequences of shifted linear systems. *Appl. Numer. Math.*, 81:105–118.
- Stein, M. L. (2014). Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics*, 8:1–19.

- Stein, M. L., Chi, Z., and Welty, L. J. (2004). Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(2):275–296.
- Stroud, J. R., Stein, M. L., and Lysen, S. (2014). Bayesian and maximum likelihood estimation for Gaussian processes on an incomplete lattice. *arXiv preprint arXiv:1402.4281*.
- Sun, Y., Li, B., and Genton, M. G. (2012). Geostatistics for large datasets. In *Advances and challenges in space-time modelling of natural events*, pages 55–77. Springer.
- Vecchia, A. V. (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 297–312.
- Vecharynski, E. and Knyazev, A. V. (2013). Absolute value preconditioning for symmetric indefinite linear systems. *SIAM J. Sci. Comput.*, 35(2):A696–A718.
- Wood, A. T. and Chan, G. (1994). Simulation of stationary Gaussian processes in  $[0, 1]^d$ . *Journal of Computational and Graphical Statistics*, 3(4):409–432.